

11-1-2011

# Autonomous Legged Hill and Stairwell Ascent

Aaron M. Johnson

*University of Pennsylvania*, aaronjoh@seas.upenn.edu

Matthew T. Hale

*University of Pennsylvania*, matthale@seas.upenn.edu

G. C. Haynes

*University of Pennsylvania*

D. E. Koditschek

*University of Pennsylvania*

---

Suggested Citation:

Aaron M. Johnson, Matthew T. Hale, G. C. Haynes, and D. E. Koditschek. "Autonomous Legged Hill and Stairwell Ascent." *IEEE International Symposium on Safety, Security, and Rescue Robotics*, November 2011, Kyoto, Japan, pp 134-142

©2011 IEEE. Personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution to servers or lists, or to reuse any copyrighted component of this work in other works must be obtained from the IEEE.

This paper is posted at ScholarlyCommons. [http://repository.upenn.edu/ese\\_papers/601](http://repository.upenn.edu/ese_papers/601)

For more information, please contact [repository@pobox.upenn.edu](mailto:repository@pobox.upenn.edu).

| Report Documentation Page  |                                    |                                     |  | Form Approved<br>OMB No. 0704-0188                  |                                    |
|--|------------------------------------|-------------------------------------|--|---|------------------------------------|
| Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.   |                                    |                                     |  |   |                                    |
| 1. REPORT DATE<br><b>01 NOV 2011</b>   |                                    | 2. REPORT TYPE                      |  | 3. DATES COVERED<br><b>00-00-2011 to 00-00-2011</b> |                                    |
| 4. TITLE AND SUBTITLE<br><b>Autonomous Legged Hill And Stairwell Ascent</b>  |                                    |                                     |  | 5a. CONTRACT NUMBER                                 |                                    |
|  |                                    |                                     |  | 5b. GRANT NUMBER                                    |                                    |
|  |                                    |                                     |  | 5c. PROGRAM ELEMENT NUMBER                          |                                    |
| 6. AUTHOR(S)   |                                    |                                     |  | 5d. PROJECT NUMBER                                  |                                    |
|  |                                    |                                     |  | 5e. TASK NUMBER                                     |                                    |
|  |                                    |                                     |  | 5f. WORK UNIT NUMBER                                |                                    |
| 7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)<br><b>University of Pennsylvania,Electrical &amp; Systems Engineering,200 S. 33rd St,Philadelphia,PA,19104</b>  |                                    |                                     |  | 8. PERFORMING ORGANIZATION<br>REPORT NUMBER         |                                    |
| 9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)  |                                    |                                     |  | 10. SPONSOR/MONITOR'S ACRONYM(S)                    |                                    |
|  |                                    |                                     |  | 11. SPONSOR/MONITOR'S REPORT<br>NUMBER(S)           |                                    |
| 12. DISTRIBUTION/AVAILABILITY STATEMENT<br><b>Approved for public release; distribution unlimited</b>  |                                    |                                     |  |   |                                    |
| 13. SUPPLEMENTARY NOTES<br><b>IEEE International Symposium on Safety, Security, and Rescue Robotics, November 2011, Kyoto, Japan, pp 134-142, Government or Federal Purpose Rights License</b>   |                                    |                                     |  |   |                                    |
| 14. ABSTRACT<br><b>This paper documents near-autonomous negotiation of synthetic and natural climbing terrain by a rugged legged robot,achieved through sequential composition of appropriate perceptually triggered locomotion primitives. The first, simple composition achieves autonomous uphill climbs in unstructured outdoor terrain while avoiding surrounding obstacles such as trees and bushes. The second, slightly more complex composition achieves autonomous stairwell climbing in a variety of different buildings. In both cases, the intrinsic motor competence of the legged platform requires only small amounts of sensory information to yield near-complete autonomy.Both of these behaviors were developed using X-RHex, a new revision of RHex that is a laboratory on legs, allowing a style of rapid development of sensorimotor tasks with a convenience near to that of conducting experiments on a lab bench. Applications of this work include urban search and rescue as well as reconnaissance operations in which robust yet simple-to-implement autonomy allows a robot access to difficult environments with little burden to a human operator.</b> |                                    |                                     |  |   |                                    |
| 15. SUBJECT TERMS  |                                    |                                     |  |   |                                    |
| 16. SECURITY CLASSIFICATION OF:  |                                    |                                     | 17. LIMITATION OF<br>ABSTRACT<br><b>Same as<br/>Report (SAR)</b> | 18. NUMBER<br>OF PAGES<br><b>10</b>                 | 19a. NAME OF<br>RESPONSIBLE PERSON |
| a. REPORT<br><b>unclassified</b>   | b. ABSTRACT<br><b>unclassified</b> | c. THIS PAGE<br><b>unclassified</b> |  |   |                                    |

# Autonomous Legged Hill and Stairwell Ascent

Aaron M. Johnson, Matthew T. Hale, G. C. Haynes, and D. E. Koditschek

Electrical & Systems Engineering, University of Pennsylvania

200 S. 33rd St, Philadelphia, PA 19104

{aaronjoh,matthale,gchaynes,kod}@seas.upenn.edu

**Abstract** — This paper documents near-autonomous negotiation of synthetic and natural climbing terrain by a rugged legged robot, achieved through sequential composition of appropriate perceptually triggered locomotion primitives. The first, simple composition achieves autonomous uphill climbs in unstructured outdoor terrain while avoiding surrounding obstacles such as trees and bushes. The second, slightly more complex composition achieves autonomous stairwell climbing in a variety of different buildings. In both cases, the intrinsic motor competence of the legged platform requires only small amounts of sensory information to yield near-complete autonomy. Both of these behaviors were developed using X-RHex, a new revision of RHex that is a *laboratory on legs*, allowing a style of rapid development of sensorimotor tasks with a convenience near to that of conducting experiments on a lab bench. Applications of this work include urban search and rescue as well as reconnaissance operations in which robust yet simple-to-implement autonomy allows a robot access to difficult environments with little burden to a human operator.

**Keywords:** *autonomous robot, hill climbing, stair climbing, sequential composition, hexapod, self-manipulation*

## I. INTRODUCTION

We present two applications of guarded autonomy for a legged robot, allowing a perceptually and algorithmically simple platform to negotiate non-trivial indoor and outdoor environments thanks to its well designed reflex and feedback mediated controls. The term *reflex* [1] denotes a purely mechanical loop arising from the interaction of a designed, shaped body or compliant limb with some naturally occurring geometric and mechanical features of the robot’s environment. The feedback policies we use all approach the ideal (and in many cases represent a formal instantiation) of an attractor-basin selected by some state-based switching logic implementing the “prepares” relation according to the sequential composition method proposed in [2]. Thus, the phrase *algorithmically simple* refers to our robot’s sole reliance on hybrid composition of online controllers to achieve guarded autonomy.

We focus on two scenarios generally acknowledged to hold great importance yet still pose considerable difficulty for existing man-portable mobile robots: the autonomous climbing of cluttered, forested hillsides [3] (Figure 1); and multi-flight stairwells in indoor settings [4] (Figure 2). In each scenario, we posit a very simple, deterministic world model and an equally simple deterministic perceptual model, along with a family of feedback controllers selected using (a sometimes slightly relaxed form of) sequential composition [2] in a manner that seems intuitively sufficient to achieve the specified navigation task. We justify that intuition by reporting



Fig. 1: The X-RHex robot on a forested hill.

extensive experimental results. Motivated by that empirical success, future versions of this work will focus with greater analytical precision on the relationships between the formal task representation, world model, algorithmic correctness and the perceptual endowment required to support it.

This new advance of guarded autonomy represents an appropriate debut for our re-engineered version of the RHex [5] platform, X-RHex [6], whose slightly greater power density and significantly more flexible sensor interface and software API enable this physical implementation of the commanded behavior that would not likely be possible for its predecessor.

### A. Motivation

For urban search and rescue (USAR) and intelligence, surveillance, reconnaissance (ISR) operations, the ability of a robot to autonomously navigate both indoor and outdoor environments provides great utility to remote operators [7]. As a typical application of our first task, autonomous ascent of a forested hillside, a robot might climb a hill to reach potential vantage points or to act as a radio relay antenna, potentially important for ISR operations as the behavior does not rely upon GPS signals. This work was motivated by preliminary tests of such a mission in the Mojave desert revealing that with relatively simple gradient-style control (see Section III-A.1) the robot climbed to the top of a small rocky hill (Figure 7). The robot encountered infrequent entrapments in the “shadow” of insurmountable but potentially easily avoidable big obstacles, thereby suggesting the need for the slightly more advanced autonomy presented in this paper. As a typical application of our second task, autonomous stairwell ascent,

a robot endowed with this capability could reach otherwise inaccessible portions of an abandoned or damaged building environment. In both settings automating the robot’s mobility to the extent of removing the detailed challenges of the local terrain from the burden on human attention (as well as on the communications channel bandwidth) further promotes its use in communications-denied or -limited environments [8].

## B. Contributions

To the best of our knowledge, no previous authors have documented the completely autonomous ascent of naturally wooded or rocky hillsides, nor of general multi-floor stairwells—much less achieving both tasks with the same robot platform. The primary contribution of this paper is our partial success in doing so on a variety of terrains (and building interior styles), documented in the data tables of Section IV. Past work in hill climbing has reported either simulation results only [9] or achieved success only through recourse to detailed terrain labeling and mapping so as to preclude failure by entrapment from minor obstacles [3, 10]. Prior work on general autonomous stairwell negotiation also has been largely focused on simulation studies [11], with almost all empirical work confined to the traversal of a single flight and yaw control on the stairs (summarized in [4]). The only prior report we have found documenting empirical work over multiple flights of stairs assumed a very specific, simple landing geometry [12]; we intentionally target a great diversity.

More broadly, we believe this work makes a secondary contribution to the literature by exploring the benefit of a greatly abstracted world model (and the greatly simplified perception required to support it) when a simple task is assigned to a mechanically competent platform. Navigation behaviors have been dominated over the last decade by interest in learning [13, 14] and, more specifically, applications of Bayesian map-building [15]. Even in their more relaxed topological representations [16], such methods are committed to repeated measurements as a necessary means of discovery, even when used on legged platforms [11]. However, the dynamics of locomotion inherent to dexterous machines such as the legged robot used in this work complicate considerably the task of accurately estimating state or building a world model [17, 18]. Here, contrarily, given the very much more narrow requirements of the task at hand, we are able to presume a priori knowledge of a “perfect” model (Section II-B). Its accuracy of course derives from its utter simplicity, inviting in turn very simple sensors. The gross discrepancies of this model with respect to the real geometry and mechanics of the environment are successfully abstracted by the mechanical reflexes of the platform.

## II. ROBOT AND TASK

### A. The Robot

1) *X-RHex, A Laboratory on Legs*: In this section we introduce the new experimental platform used in this paper, X-RHex [6]. Shown in Figures 1 and 2, X-RHex has about the the same footprint and weight as Research RHex [5], but only half the body height. Its motors are 2.5 times stronger,



Fig. 2: The X-RHex robot on a set of stairs with laser scanner, IMU, wireless repeater, and handle payloads.

making them useful for both climbing hills and stairs. The robot can slot-load up to two batteries, each of which lasts roughly 1.5 times the original RHex battery, enabling longer experimental runs. A full report on the platform and a detailed comparison to past RHex robots can be found in [6].

One significant advantage of the new platform, and a design extension beyond prior RHex platforms, is the introduction of a payload system on the top of the robot, the space for which is afforded by the robot’s thinner profile. The system consists of a standardized mechanical mount, and a set of electrical connectors to interface the payloads with on-board electronics. With swappable payloads, the robot functions as a *laboratory on legs* and supports an open-ended variety of experiment-specific sensory and computational payloads. In these experiments we use a laser scanner<sup>1</sup> and IMU<sup>2</sup>, as well as an additional wireless communications payload and a pair of carry handles.

A second major advance over prior RHex platforms is the new “Dynamism” [6] development environment, providing a lightweight interface to store and retrieve data, either from other functions or processes on the robot or from other computers on the network. For example, the locomotion primitives we use in these experiments are all coded in compiled executables on the robot, whereas the sensor-based behaviors developed in this paper have been coded in a scripting language (Python or MATLAB) on a laptop client for simplicity. While all these behaviors could be coded directly on the robot, the use of this network abstraction layer has greatly sped up behavior development, though occasional network glitches caused some problems in the experiments (as we document below).

2) *Abstract Robot Model*: For purposes of task specification and modeling, we assume the robot’s standard gait (alternating tripod [5]) over the standard terrain encountered (as modeled in the next section) reduces to the target dynamics (or “template”

<sup>1</sup>Hokuyo URG-04LX-F01, <http://www.hokuyo-aut.jp/>, an indoor unit that was used outdoors but not in direct sunlight.

<sup>2</sup>Microstrain 3DM-GX2, <http://www.microstrain.com/>




Fig. 3: Position tracks of several runs up the same hill with and without automatic uphill steering (blue solid and red dashed lines, respectively), with various starting angle. Uphill is positive Y direction (North), axis are in meters.

[19]) of a horizontal plane kinematic unicycle [20],

$$\dot{x} = v \sin(\theta) \quad (1)$$

$$\dot{y} = v \cos(\theta) \quad (2)$$

$$\dot{\theta} = s \quad (3)$$

controlled by a velocity ( $v$ ) and steering ( $s$ ) command. Physically, when RHex climbs at an angle to an uphill direction, gravity will naturally yaw the robot downhill. This can be seen in the red curves of Figure 3, which show a number of trials of the robot walking on a hill from various initial headings and no steering command ( $s = 0$ ). These data suggest a more realistic model for heading dynamics would take the form

$$\dot{\theta} = s + \delta \sin(\alpha) \sin(\phi) \quad (4)$$

where  $\phi$  is the local vertical slope and  $\alpha$  is the yaw angle of the robot's heading relative to the direction of the slope<sup>3</sup>.

### B. The World Model

We now introduce the very simple “grade” model of a terrain that will abstract away almost all the physical properties of the hills and stairs to provide a uniform view of the robot's task within its environment. This abstraction is only appropriate on a platform such as RHex whose normal walking gait can safely handle small obstacles (rocks, twigs, etc).

1) *The Grade Terrain Model:* A *terrain* is specified by some (unknown) height function,  $\eta \in C^\infty[\mathbb{R}^2, \mathbb{R}]$ . Not only is  $\eta$  unknown, but we assume it is not a metrically full scale accurate copy of the literal terrain, rather to be imagined as sufficiently “smoothed” and thus absent of spatial frequencies much below the robot's bodylength. Its (also unknown) gradient,

$$D\eta(\mathbf{x}) = \gamma(\mathbf{x}) \cdot \widehat{D\eta}(\mathbf{x}); \quad \gamma(\mathbf{x}) := \|D\eta(\mathbf{x})\|, \quad (5)$$

we write in polar form as the product of the *grade*,  $\gamma$ , and *steepest ascent* unit field,  $\widehat{D\eta}$ . The set of *obstacles* is given by excessively steep grades

$$\mathcal{O} := \{\mathbf{x} \in \mathbb{R}^2 : \gamma(\mathbf{x}) \geq G_0\},$$

<sup>3</sup> We ascribe these effective yaw perturbation forces to the overall consequences of the “downhill” legs taking more of the robot's weight and thus lagging behind the “uphill” legs. In particular, note that the magnitude of the effect gets worse the farther the robot turns (modeled by the  $\sin(\alpha)$  term).

where  $G_0$  is a lower bound on the grades above which the alternating tripod gait will not successfully propel the machine in a manner well modeled by the unicycle plant introduced above in equations (1)-(4).

We conjecture (but do not attempt to rigorously establish in this paper) that the sequential composition methods to be introduced in the next section can be proven correct under the assumption that the terrain is “simple”: i.e., that the obstacle set of excessively high grades comprises a disjoint union of “suitably” separated (defined as a gap wide enough to fit through a proximity-distance-sensor thickened disk containing the robot's horizontal plane body) convex compact shapes. Under these circumstances, the obstacle-free planar surface on which the robot navigates is a topological sphere world in the sense of [21]. We assume throughout the rest of this paper that the actual terrain has this property (and report only on the empirical aspects of the resulting implementation).

2) *Hill and Stair Models and Climbing Tasks:* A *hill* is any simple terrain. We define the *hill climbing task* as requiring that the robot locomote from any initial position and orientation to some local maximum of the height function  $\eta$ .

In contrast, we define a stairwell to be a piecewise constant terrain (each constant component called a *landing*) with obstacle boundaries (walls, cliffs) including a distinguished subset called a *stair* that connects the landings. We will define a stair purely in terms of its perceptual features as detailed below in Section II-C.6. Unlike other “excessively steep” terrain, a stair can be ascended by recourse to a different gait (described in Section III-B). The *stairwell climbing task* requires that the robot locomote from any initial position and orientation in a stairwell to some landing with no (upward) “stair” boundaries.

### C. Sensor Models

In this section we posit a simple set of abstract sensor models and briefly relate how they are realized (of course, actually, merely approximated) in our physical hardware. First, we introduce a vestibular sensor relying only on a conventional IMU output, and then a succession of exteroceptive sensors that can be realized through use of a LIDAR hardware unit mounted on a legged robot.

1) *Gravitational Gradient Sensor:* Given the orientation of a robot's body from an IMU (in terms of a coordinate system  $\mathbf{x}$ ,  $\mathbf{y}$ , and  $\mathbf{z}$ ), the calculation of the instantaneous uphill direction is similar to computations proposed in the prior stairwell experiments [12]. We compute the rotation  $\alpha$  about  $\mathbf{z}$  between  $\mathbf{x}$  and  $\widehat{D\eta}$ , given the direction of gravity,  $\mathbf{g}$ , as follows

$$\alpha := \arccos(\mathbf{x} \cdot [(\mathbf{g} \times \mathbf{z}) \times \mathbf{z}]) \quad (6)$$

2) *Excessive Grade Sensor:* The excessive grade sensor is an abstract depth map,

$$\sigma_E : \mathbb{R}^2 \times \mathbb{S}^1 \times [-P, P] \times [-A, A] \rightarrow [0, R]$$

that returns from each position and orientation in the plane,  $(x, y, \theta) \in \mathbb{R}^2 \times \mathbb{S}^1$ , body pitch,  $\psi \in [-P, P]$ , and view direction,  $\lambda \in [-A, A]$ , a distance,  $\rho \in [0, R]$ , to the nearest excessive grade. In our implementation, we use the output






Fig. 4: The pitch wiggle behavior, with middle legs removed for clarity.

from a fixed LIDAR unit to realize this depth map. The arc extends roughly  $\pm A$  (where  $A = 120^\circ$ ) off center. The distance profile corresponds to the first depth at which the LIDAR unit records a return. For the chosen fixed placement of this unit, our robot will interpret as an obstacle anything (tree, rock, slope increase, wall) that rises more than 25cm over a 1m run above the existing slope — hence, abstractly, this sensor is indeed responding to an excessively steep grade,  $\gamma > G_0$ , corresponding to the terrain model above. The LIDAR unit cannot “see” beyond a distance of  $R := 4m$ , to which the “infinite” reading of its maximum depth scale is calibrated. The laser scanner plane is at a height such that any obstacle that it cannot see is assumed to be surmountable and any obstacle that it can see is assumed to be insurmountable.

3) *Gap Sensor* : The gap sensor is an abstract map,

$$\sigma_G : \mathbb{R}^2 \times \mathbb{S}^1 \rightarrow [-A, A]$$

that returns for each position and orientation at which the robot is pointing, the center,  $\sigma_G(x, y, \theta) = \xi$  of an arc segment  $[\xi - S, \xi + S] \subset [-A, A]$ , a “window” within which the interval depth is maximized

$$\xi := \operatorname{argmax}_{\tau \in [-A+S, A-S]} I[\tau, S]$$

where the “minimum interval depth” is taken to be

$$I_M[\alpha, \beta] := \min_{\alpha - \beta < \lambda < \alpha + \beta} \sigma_E(x, y, \theta, 0, \lambda)$$

resulting in the sensor  $\sigma_{G_M}$  (to be used in the hill climbing application) whereas the “weighted mean” depth is taken to be a distorted mean

$$I_W[\alpha, \beta] := \min_{\alpha - \beta < \lambda < \alpha + \beta} \frac{\sigma_E(x, y, \theta, 0, \lambda)}{(1 - K)\cos^6(\lambda - \alpha) + K}$$

resulting in the sensor  $\sigma_{G_W}$  (to be used in the stairwell climbing application).

4) *Plane Sensor*: The plane sensor,  $\sigma_P$  is defined as,

$$\sigma_P(x, y, \theta) := \{(\rho, \lambda, \psi) : \rho = \sigma_E(x, y, \theta, \psi, \lambda)\}$$

and is implemented by running an excessive grade sensor at each pitch angle achieved via the “pitch wiggle” self-manipulation.

The pitch wiggle is a sensorimotor routine which uses the planar LIDAR to measure ranges in many planes. With a LIDAR unit such as this positioned horizontally with respect to the ground, a staircase for example will appear similar to a wall. Unlike many robots that attach a LIDAR unit to a

motorized tilting mechanism, we use RHex’s natural ability to self-manipulate to a variety of angles in order to sweep the LIDAR’s sensing plane. This maneuver produces a large variation in body pitch (either up or down) with minimal internal forces or toe slip<sup>4</sup>, and is depicted in Figure 4. A more precise treatment of robot self manipulation RHex will be presented in an upcoming paper [22].

5) *Cliff Sensor*: The cliff sensor,  $\sigma_C$ , is a simple threshold on the ranges returned from the pitch wiggle,  $\sigma_P < C$ , where  $\sigma_P$  is pitched downward and  $C$  is a calibrated range verifying that the ground is where we expect it to be.

6) *Stair Sensor*: The stair sensor  $\sigma_S$  is the composition  $\sigma_F \circ \sigma_P$ , where  $\sigma_P$  is pitched around an upward center and  $\sigma_F$  tries to find the most “stepped” direction<sup>5</sup>. Note that this stairwell detector relies on the assumption that all walls and obstacles are infinitely high, that is there are no other “stepped” features on the landings.

### III. DESCRIPTION OF AUTONOMOUS BEHAVIORS

#### A. Autonomous Hill Climbing

1) *Steepest Ascent Controller*: For this controller, the fore-aft velocity command is set to a grade-selected constant

$$v := v_0 \chi(\gamma/G_0), \quad (7)$$

where  $\chi \in C^\infty[0, 1]$  fixes 0, i.e.,  $\chi(0) = 0$ , so that locomotion ceases when the grade is level<sup>6</sup>. The key task of uphill gradient following is achieved by introducing a servo term<sup>7</sup>

$$s := k_s \alpha, \quad (8)$$

where  $\alpha$  is the angular in heading difference from steepest ascent (see equation (6)). Assuming that  $k_s > \delta \sin \phi$ , this stabilizes the robot’s heading angle around the true steepest ascent direction. The solid blue lines in Figure 3 record experiments using the robot with this controller, on the same hill as before, with a wide variety of initial angles (including one facing directly downhill) forced successfully toward the heading of steepest ascent<sup>8</sup>.

It is intuitively clear (but, adhering to the intended informal scope of the present paper, we will not rigorously prove) that this controller yields a closed loop system when applied to the unicycle plant model, equations (1)-(4), respecting which the terrain height,  $\eta$ , represents a LaSalle function

<sup>4</sup>This pitch change can be easily derived from the geometry of the C-shaped legs.

<sup>5</sup>The details of  $\sigma_F$  are not important, but it is basically a well filtered edge detector that is tuned to be sensitive to the magnitudes and frequencies of typical stairwells as seen by a laser scanner sweeping through a range of pitches. Importantly, we have set these parameters once and for all, and they remain invariant throughout all the different stairwells of Table II.

<sup>6</sup>The leg offset parameter of the gait is also automatically modified as the robot runs to match the grade as in [23].

<sup>7</sup>Here and elsewhere in this paper  $s$  is filtered to smooth out noise and prevent a large instantaneous change in steering command, which also gives these controllers some effective (and pragmatically welcome) hysteresis.

<sup>8</sup>In general, servoing on a direction other than steepest ascent would require additive feedforward cancellation of the yaw perturbation effect. While estimates of  $\phi$ , the pitch, can be readily achieved for this, the effective yaw perturbation gain,  $\delta$ , varies significantly across terrain types so that some online identification or adaptive control technique would be required to achieve reliable performance.




Fig. 5: A cylindrical projection of a patch of the world model with controller superimposed for the hill steering behavior.

[24]. Consequently, the resulting attracting set consists of the maxima of  $\eta$  together with the subset of the obstacle boundary whose tangent space is exactly normal to the heading direction. In other words, the robot will keep climbing along the direction of steepest ascent until it has either halted at a local maximum or has “crashed into” an obstacle defined by the excessive grade condition,  $\gamma \geq G_0$ . The basins of attraction around each of these components of the forward limit set are typically complicated open regions in the set of position and heading initial conditions determined by the precise details of the terrain,  $\eta$ . Of course, the central point of this paper is that the robot does not need to know nor learn anything about them. It is sufficient to know that this gradient ascent controller creates a multi-stable family of attractors, some of which correspond to the completion of the climbing task, while others, the “obstacle crash encounters,” represent undesired events we will now place in sequential composition with avoidance controllers.

2) *Obstacle Avoidance Controller:* We can replace the controller of equation (8) with a “gap seeking” steering command,

$$s := k_s \sigma_{G_M}(x, y, \theta), \quad (9)$$

and then use the assumption that obstacles are convex to guarantee that  $\sigma_{G_M}$  must eventually return to  $0^\circ$  at which point the “prepares” relation is triggered again (i.e. the machine is in a new steepest ascent basin) and the uphill steering behavior takes over. Note as another consequence of the obstacle convexity assumption that this obstacle avoidance behavior will still make progress in terms of our LaSalle function,  $\eta$ , since the robot will never turn downhill.

3) *Combined Behavior:* Figure 5 summarizes the behavior, and shows a metrically deformed, topologically accurate rendering in a  $\theta \equiv 90^\circ$  projected slice of the robot’s  $(x, y, \theta)$  space of the attractor-basin family arising from sequentially composing the steepest ascent controller (equations 7 and 8) and the obstacle avoidance controller (equation 9) in the face of a hill,  $\eta$ . The black arrows depict the typical controller-induced vector field. The top purple line represents a component of the attracting terrain “goal” maxima. The red disks represent excessive-grade obstacles and they each induce an attracting “crash” maxima on their boundary. The pre-image

basin of the “goal” attractor is depicted in yellow while the basin of the “crash” attractors is depicted in green. The domain of the obstacle avoidance controller is represented in the blue-dotted region located just proximal to the red obstacles.

### B. Autonomous Stairwell Ascent

Because of the additional perceptual and motor activity associated with finding and negotiating stairs, the autonomous stairwell ascent behavior has greater complexity than the hill climbing just presented. Although we again address the overall task through the systematic construction of pre-image backchaining [25], our reliance upon reflexes implies that not all the action steps will admit well-defined attractors and basins as required for the very robust and formally more powerful variant of sequential composition [2]. We report here on the presently functioning constituents of this behavior and leave for future work their formal reconciliation into that more powerful (but restrictive) framework.

1) *The Stair Climbing Behavior:* RHex robots have been climbing single-flight stairs for nearly a decade since Buehler’s group first developed the appropriate gait [26] and they perform quite reliably on a variety of typical human-scale staircases. This capability owes much to the reflex yaw stabilization conferred by in-phase contra-lateral legs (providing a wide base of support on each successive stair) along with the metachronal gait that engages the circular legs just in time to place the body weight quasi-statically on the tread of the stair. The reflexes arising from this gait ensure that RHex-style legged platforms ascend stairs in open-loop as if in the presence of the perceptually active steepest ascent stabilizing controller (equation 8) on hills<sup>9</sup>.

When climbing a single-flight stairwell, we can consider backchaining through a (still to be defined) fuller palette of controllers in the spirit of [2]. Our goal set is the top of the final flight of stairs. Thus we need one controller that can direct the robot to that goal set from say just past the top step of the final flight. This stair exit controller is simply a few walking steps triggered by the robot body pitch (as done for a different robot in [12]).

We have already described a controller that climbs the stairs and has a reflex funnel from a domain of any individual step (say the first one) to any higher step (such as the last step). In order to enter the domain of the stair climbing controller, the robot uses a transition from walking to stair climbing that has also been shown to be reliable when the robot is walking towards the start of a stairwell [28]<sup>10</sup>.

2) *Landing Behavior:* Now we have successfully backchained a series of dynamical controllers that can guarantee the robot can get from a lower landing in front of a stairwell to the next landing just past the top of a stairwell. A new sequence of controllers is needed to get the robot from the goal set of one flight to the starting domain of the next.

<sup>9</sup>And, as it turns out, at least as reliably in this task open-loop mode as any tracked robots under feedback control since the latter must place their weight on the nose of the stair for each step [27], which is contrary to the way stairs are intended to be used.

<sup>10</sup>Now using the virtual contact sensor of [29] instead of mere controller error to trigger the same transition more reliably.

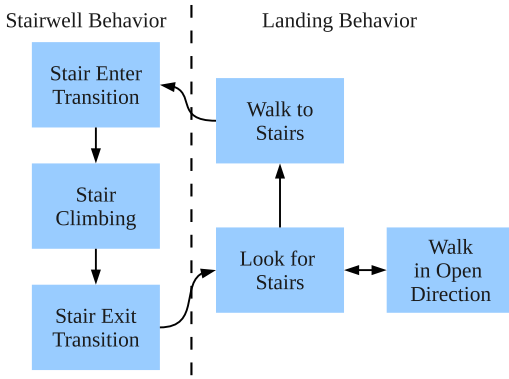


Fig. 6: Flow chart describing autonomous stair climbing.

First this controller uses the stair sensor  $\sigma_S$  (Section II-C.6), and will perform an open-loop<sup>11</sup> walk to the next stairs if seen. If the robot is not sure that there is a staircase in view it then moves on to explore the landing. The robot will pick a direction that is “most open” by using the gap sensor  $\sigma_G$  (Section II-C.3). The robot then performs an open-loop walk in that direction for about a meter and then stops to look for stairs again. If no suitably open direction is found (if  $\sigma_E(\sigma_{G_W}) < 1m$ ) the robot simply turns  $90^\circ$ . In order to guarantee that the robot will not fall off any cliffs (such as the down stairs direction), the robot checks the cliff sensor  $\sigma_C$  (Section II-C.5) to verify that there is navigable terrain in front of it.

This would be a very slow way to explore an entire building, and there are no deterministic guarantees that the robot must find the next flight of stairs. However given the fact that landings are generally metrically small and not maze-like, this controller will in reality find the subsequent stairwell with a very high probability.

With this new “explore a landing” behavior the sequential composition backchaining is completed and produces a cyclic path through controllers until the robot reaches the top of a stairwell. Figure 6 summarizes the entire behavior.

#### IV. EXPERIMENTAL RESULTS

##### A. Autonomous Hill Climbing

To test the autonomous hill climbing behavior, a dataset was collected on eight different hills located in two nearby parks, as summarized in Table I. Overall, the robot climbed almost half a kilometer of hilly terrain while avoiding 42 obstacles (21 trees, 10 bushes, 10 logs, 1 human) and hitting only 5 obstacles. In other words, the steepest ascent controller, with no obstacle avoidance controller, would otherwise have hit and likely become entrapped by an additional 42 obstacles that the robot was able to avoid in these tests. Larger obstacles, such as very wide trees or downed logs, took significant amounts of time to follow around and the follow time for any obstacle of at least 3s has been added up and reported.

Four of the eight hills had a relatively consistent grade, and their hill slope is listed as an average value, while the

other hills were more uneven and their hill slope is listed as a range. All of these latter runs encountered some terrain that was nearly flat (hence, the robot, achieving a local maximum, formally accomplished the task at hand in these instances), however we disabled the summit detection to allow the robot to possibly cross these intermediate plateaus so as to keep accumulating climbing statistics. Runs 2 and 6 ended when the robot reached such a *plateau* because the robot ended up drifting in yaw back toward its initial placement. Trial 8 ended in a *summit* and trial 4 ended when it reached the *edge* of the park (an artificial boundary that violates the world model). The other four trials were stopped when the robot incurred a fault after the reported distance had been covered. Three trials were stopped with either a *flip* (robot inverted when it hit a low log) or a *peg* (robot stuck on a small but rigid branch), and one trial was stopped due to the *hill* getting too steep. All are failures in the assumption that any obstacle that the sensor cannot see is surmountable (Section II-C.2), and could be fixed with either better sensing or better locomotion primitives.

##### B. Autonomous Stairwell Ascent

To test the autonomous stairwell climbing behavior we ran it on 10 of the many different stairwells in 4 nearby buildings<sup>12</sup>, as Table II summarizes<sup>13</sup>. We distinguish *behavior faults* (arising from inadequacies in either the algorithm or the sensorimotor capabilities that subserve it) from *robot faults* (failures due to mechanical or electronic unreliability). Only four of the stairwells met the requirements of our world model. Specifically, they exhibited solid, detectable walls and no significant stepped features on the landings. These first four runs listed yielded cumulatively only two behavioral problems throughout the ascent of 103 steps accumulated over 15 flights. Both of these behavioral failures were incurred by faults arising from a false negative in the *cliff* sensor.

The rest of the stairwells violated our world model assumptions, but the robot was still able to climb with varying degrees of success. Their most significant discrepancy was their glass or mesh *walls* that the laser scanner can see through but are actually obstacles to the robot, leading to either a direct or indirect collision that in turn precipitated faults requiring operator intervention. The mesh stairs (the robot’s worst performance) as well as one stairwell with boxy heaters both look very “stepped” (as described in Section II-C.6), and lead to several *stair* faults — false positives in the stair sensor. Other behavioral faults included problems in the *transition* between walking in stair climbing, usually due to either bad starting angle or premature transition.

In addition to behavioral faults, there were 17 robot faults over all 61 flights. The majority of these arise from a *leg* failing to respond or from a robot *reset* - both due to a known

<sup>12</sup>There were three additional stairwells that were attempted but on which the robot made no progress due to their having either open risers or glossy painted risers that the laser scanner could not see well if at all. This is a limitation of the sensor and these stairwells are not reported.

<sup>13</sup>Naturally every stairwell is unique, and even within a stairwell the rise, run, width, steps per flight, landing size, style, and wall type can vary significantly. Listed here are “typical” values for a given stairwell that attempt to convey some of these differences without providing full blueprints.

<sup>11</sup>In the current version, the open-loop walking parameters are floor material specific. However this hand-crafted step could readily be eliminated in favor of some form of terrain identification or visual servoing.



| # | Description          | Direct Distance | Hill Slope | Runtime | Obstacles Avoided | Follow Time | Faults    | Finish  |
|---|----------------------|-----------------|------------|---------|-------------------|-------------|-----------|---------|
| 1 | Steep, Rocky Hill    | 14 m            | 17°        | 1:35    | 3                 | 3 s         | Hill      | Fault   |
| 2 | Medium Grassy A      | 26 m            | 0 – 13°    | 1:23    | 1                 | 4 s         | -         | Plateau |
| 3 | Steep, Sparse Forest | 48 m            | 14°        | 3:26    | 2                 | 0 s         | Peg       | Fault   |
| 4 | Sparse Grassy A      | 51 m            | 0 – 8°     | 4:54    | 2                 | 0 s         | -         | Edge    |
| 5 | Medium Forest A      | 57 m            | 11°        | 4:20    | 10                | 25 s        | Flip      | Fault   |
| 6 | Medium Grassy B      | 62 m            | 0 – 8°     | 5:42    | 7                 | 4 s         | -         | Plateau |
| 7 | Medium Forest B      | 69 m            | 9°         | 4:37    | 15                | 13 s        | Peg, Flip | Fault   |
| 8 | Sparse Grassy B      | 124 m           | 0 – 9°     | 6:55    | 2                 | 0 s         | -         | Summit  |

TABLE I: Eight outdoor hill climbing behavior trials including 42 successfully avoided obstacles over almost half a kilometer of climbing with only 5 faults.

| #  | Wall  | Rise   | Run    | Landing  | Landing Size | Flights | Stairs | Time    | Scans | Behavior Faults         | Robot Faults   |
|----|-------|--------|--------|----------|--------------|---------|--------|---------|-------|-------------------------|----------------|
| 1  | -     | 15.3cm | 28.0cm | Straight | 189x150cm    | 2       | 11     | 2:23    | 2     | -                       | 1 IMU          |
| 2  | -     | 15.3cm | 28.0cm | Straight | 327x150cm    | 2       | 11     | 2:55    | 3     | -                       | -              |
| 3  | -     | 17.4cm | 27.9cm | Mixed    | 494x321cm    | 2       | 21     | 14:29   | 17    | 2 Cliff                 | 1 LIDAR, 1 Leg |
| 4  | -     | 18.2cm | 26.3cm | U-Left   | 486x222cm    | 7       | 60     | 32:53   | 30    | -                       | 1 LIDAR        |
| 5  | Glass | 17.4cm | 29.6cm | Straight | 192x143cm    | 2       | 27     | 6:48    | 7     | 1 Stair                 | 1 Reset        |
| 6  | Glass | 16.7cm | 26.9cm | Mixed    | 256x277cm    | 3       | 25     | 23:27   | 24    | 2 Cliff, 1 Tran         | -              |
| 7  | Glass | 16.2cm | 28.5cm | U-Left   | 471x252cm    | 10      | 111    | 1:05:19 | 67    | 1 Stair, 1 Wall         | 1 Leg          |
| 8  | Glass | 17.3cm | 27.2cm | U-Left   | 349x156cm    | 10      | 112    | 40:49   | 42    | 2 Tran, 1 Cliff, 1 Wall | 2 Legs         |
| 9  | Mesh  | 17.3cm | 27.2cm | Mixed    | 293x137cm    | 11      | 112    | 1:26:59 | 89    | 8 Various               | 7 Various      |
| 10 | Boxy  | 17.5cm | 26.0cm | U-Left   | 228x122cm    | 12      | 181    | 1:11:41 | 71    | 2 Stair, 1 Tran         | 1 Leg          |

TABLE II: Ten indoor stairwell climbing behavior trials covering 671 stairs in 61 flights with a total of 23 behavioral problems.

power regulator issue scheduled for correction in the next X-RHex hardware update. Additionally, there were several sensor faults, where either the *IMU* or *LIDAR* would stop responding due to the regulator issue or a hardware, firmware, or driver problem. Finally, there were 3 faults where it appeared that some *network* corruption caused the robot to miss a command (such as stop walking) and led to a crash of some sort.

Overall the behavior was able to climb a total of 671 stairs in 61 flights while encountering only 23 behavioral faults in almost 6 hours of testing. In almost every stairwell, there were many other incidents that could be considered faults (such as a leg hitting a wall, open loop walking leaving the robot at the wrong angle, false-positive on cliff detection, etc) but the robust reflexes and reactive behaviors prevented these from requiring a human intervention.

## V. CONCLUSIONS AND FUTURE WORK

We have presented a rudimentary form of guarded autonomous locomotion whose empirically demonstrated robustness in unstructured natural and synthetic environments rests upon the underlying motor competencies of the host platform, stitched together with very simple perceptually triggered switches whose deployment is arranged in a manner idealized by the formal notion of sequential composition [2]. Although a formal argument about the correctness of the idealized hybrid controller placed within the idealized world model lies beyond the scope of the paper, we have tried to sketch along the way how we imagine such proofs to be eminently achievable.

In a more practical vein, we are convinced that several further modest extensions and improvements in the behaviors presented here would considerably close the gap to full autonomy still revealed by the tables, thereby conferring

true applications-worthy utility upon the X-RHex platform. For example, the stair climbing behavior can be endowed with descent capability (as in [30] via [31]), as well more deliberative obstacle avoidance (as in [32]). The hill climbing behavior could use grade to cue a greater diversity of better hill climbing gaits [33] in order to climb hills as steep as 45°.

We are planning on expanding the range of hills and stairs that the behavior will be tested on, starting back in the Mojave desert as seen in Figure 7. Our preliminary tests have shown that the steepest ascent controller of Section III-A.1 works well even on very rough terrain, and we suspect that the obstacle avoidance behavior will allow the robot to avoid hitting the larger rocks and bushes.

Harking back to the initial theme of sensor minimality, we also suspect that both of these behaviors could be completed using no exteroceptive sensors at all. Instead the robot would rely on proprioceptive sensors and use the legs to “feel” obstacles. For stair climbing, we could use a “virtual contact sensor” to feel the walls and a “missing ground” sensor as a cliff detector [29]. For autonomous hill climbing, the gravitational gradient sensor could be replaced with a proprioceptive (instead of vestibular) sensor [34].

Finally, we conjecture that legs matter: that these same behavioral notions applied to a simpler platform (such as a tracked or wheeled robot) would not perform as well, most notably because the robot would get stuck on obstacles that the mechanical reflexes allow us to ignore. RHex does not need nor want perfect sensing and we aim to wring the most autonomy possible by appeal to the simplest most robust sensor available. We seek to tell the robot not to hit a rock, or direct it generally towards the stairs. After that, we hope to



Fig. 7: The X-RHex robot on a rocky desert hill.

let the “mechanical intelligence” take over the details.

#### ACKNOWLEDGEMENTS

This work has been supported in part by a DARPA Max-Mobility Seedling grant and the Army Research Laboratory under Cooperative Agreement Number W911NF-10-2-0016, and in part by the National Science Foundation under NSF CDI-II Number 1028237, with additional support from an Intelligence Community Postdoctoral Fellowship held by the third author. The authors would like to thank David Hallac, Praveer Nidamaluri, Ryan Knopf, and the rest of the lab for their help while writing this paper.

#### REFERENCES

- [1] I. E. Brown and G. E. Loeb, *A reductionist approach to creating and using neuromusculoskeletal models*. Springer, 2000, p. 148.
- [2] R. R. Burridge, A. A. Rizzi, and D. E. Koditschek, “Sequential composition of dynamically dexterous robot behaviors,” *The International Journal of Robotics Research*, vol. 18, no. 6, pp. 534–555, 1999.
- [3] K. Konolige, M. Agrawal, M. Blas, R. Bolles, B. Gerkey, J. Solà, and A. Sundaresan, “Mapping, navigation, and learning for off-road traversal,” *Journal of Field Robotics*, vol. 26, no. 1, pp. 88–113, 2009.
- [4] R. Ray, B. Bepari, and S. Bhaumik, “On design and development of an intelligent mobile robotic vehicle for stair-case navigation,” in *Intelligent Autonomous Systems*, ser. Studies in Computational Intelligence, D. Pratihar and L. Jain, Eds. Springer Berlin / Heidelberg, 2010, vol. 275, pp. 87–122.
- [5] U. Saranli, M. Buehler, and D. E. Koditschek, “RHex: A simple and highly mobile hexapod robot,” *International Journal of Robotics Research*, vol. 20, no. 7, pp. 616–631, 2001.
- [6] K. C. Galloway, G. C. Haynes, B. D. Ilhan, A. M. Johnson, R. Knopf, G. Lynch, B. Plotnick, M. White, and D. E. Koditschek, “X-RHex: A highly mobile hexapodal robot for sensorimotor tasks,” University of Pennsylvania, Tech. Rep., 2010.
- [7] A. Birk and S. Carpin, “Rescue Robotics - a crucial milestone on the road to autonomous systems,” *Advanced Robotics*, vol. 20, no. 5, 2006.
- [8] R. R. Murphy, S. Tadokoro, D. Nardi, A. Jacoff, P. Fiorini, H. Choset, and A. M. Erkmén, “Search and rescue robotics,” in *Springer Handbook of Robotics*, B. Siciliano and O. Khatib, Eds. Springer Berlin Heidelberg, 2008, pp. 1151–1173.
- [9] R. Arkin and W. Gardner, “Reactive inclinometer-based mobile robot navigation,” in *Robotics and Automation, 1990. Proceedings., 1990 IEEE International Conference on*, may 1990, pp. 936–941 vol.2.
- [10] S. Thrun, M. Montemerlo, and A. Aron, “Probabilistic terrain analysis for high-speed desert driving,” in *Proceedings of Robotics: Science and Systems*, Philadelphia, USA, August 2006.
- [11] B. Morisset, R. Rusu, A. Sundaresan, K. Hauser, M. Agrawal, J. Latombe, and M. Beetz, “Leaving flatland: toward real-time 3D navigation,” in *Robotics and Automation, 2009. ICRA’09. IEEE International Conference on*. IEEE, 2009, pp. 3786–3793.
- [12] S. Steplight, G. Engal, S.-H. Jung, D. B. Walker, C. J. Taylor, and J. P. Ostrowski, “A mode-based sensor fusion approach to robotic stair-climbing,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2000, pp. 1113–1118.
- [13] J. Pippine, D. Hackett, and A. Watson, “An overview of the Defense Advanced Research Projects Agency’s Learning Locomotion program,” *The International Journal of Robotics Research*, vol. 30, no. 2, p. 141, 2011.
- [14] E. Krotkov, S. Fish, L. Jackel, B. McBride, M. Perschbacher, and J. Pippine, “The DARPA PerceptOR evaluation experiments,” *Autonomous Robots*, vol. 22, no. 1, pp. 19–35, 2007.
- [15] S. Thrun, W. Burgard, and D. Fox, *Probabilistic Robotics*. MIT Press, 2006.
- [16] A. Ranganathan, E. Menegatti, and F. Dellaert, “Bayesian inference in the space of topological maps,” *IEEE Transactions on Robotics*, vol. 22, no. 1, pp. 92–107, 2006.
- [17] S. Skaff, A. Rizzi, H. Choset, and M. Tesch, “Context Identification for Efficient Multiple-Model State Estimation of Systems With Cyclical Intermittent Dynamics,” *Robotics, IEEE Transactions on*, no. 99, pp. 1–15, 2011.
- [18] P.-C. Lin, H. Komsuoglu, and D. E. Koditschek, “A leg configuration measurement system for full body posture estimates in a hexapod robot,” *Robotics, IEEE Transactions on*, vol. 21, no. 3, pp. 411–422, 2005.
- [19] R. Full and D. Koditschek, “Templates and anchors: Neuromechanical hypotheses of legged locomotion on land,” *Journal of Experimental Biology*, vol. 202, pp. 3325–3332, 1999.
- [20] G. A. D. Lopes and D. E. Koditschek, “Visual servoing for nonholonomically constrained three degree of freedom kinematic systems,” *The International Journal of Robotics Research*, vol. 26, no. 7, pp. 715–736, 2007.
- [21] D. Koditschek and E. Rimon, “Robot navigation functions on manifolds with boundary,” *Advances in Applied Mathematics*, vol. 11, no. 4, pp. 412–442, 1990.
- [22] A. M. Johnson, G. C. Haynes, and D. E. Koditschek, “Self Manipulation and Reactive Behaviors for RHex,” 2011, in preparation.
- [23] H. Komsuoglu, D. McMordie, U. Saranli, N. Moore, M. Buehler, and D. E. Koditschek, “Proprioception based behavioral advances in a hexapod robot,” in *Proceedings of International Conference on Robotics and Automation*, Seoul, Korea, 2001.
- [24] J. P. LaSalle, *The Stability of Dynamical Systems*. Society for Industrial Mathematics, 1976.
- [25] T. Lozano-Perez, M. T. Mason, and R. H. Taylor, “Automatic synthesis of fine-motion strategies for robots,” *The International Journal of Robotics Research*, vol. 3, no. 1, pp. 3–24, 1984.
- [26] E. Moore, D. Campbell, F. Grimmering, and M. Buehler, “Reliable stair climbing in the simple hexapod ‘RHex,’” in *Robotics and Automation, 2002. Proceedings. ICRA’02. IEEE International Conference on*, vol. 3. IEEE, 2002, pp. 2222–2227.
- [27] A. Mourikis, N. Trawny, S. Roumeliotis, D. Helmick, and L. Matthies, “Autonomous stair climbing for tracked vehicles,” *The International Journal of Robotics Research*, vol. 26, no. 7, p. 737, 2007.
- [28] G. C. Haynes and A. A. Rizzi, “Gaits and gait transitions for legged robots,” in *Proceedings of the IEEE International Conference On Robotics and Automation*, Orlando, FL, USA, May 2006, pp. 1117–22.
- [29] A. M. Johnson, G. C. Haynes, and D. E. Koditschek, “Disturbance detection, identification, and recovery by gait transition in legged robots,” in *Proceedings of the IEEE/RSJ Intl. Conference on Intelligent Robots and Systems*, Taipei, Taiwan, October 2010.
- [30] J. Hesck, G. Mariottini, and S. Roumeliotis, “Descending-stair detection, approach, and traversal with an autonomous tracked vehicle,” in *Intelligent Robots and Systems (IROS), 2010 IEEE/RSJ International Conference on*, oct. 2010, pp. 5525–5531.
- [31] D. Campbell and M. Buehler, “Stair descent in the simple hexapod RHex,” in *Robotics and Automation, 2003. Proceedings. ICRA’03. IEEE International Conference on*, vol. 1. IEEE, 2003, pp. 1380–1385.
- [32] A. Kalantari, E. Mihankhah, and S. Moosavian, “Safe autonomous stair climbing for a tracked mobile robot using a kinematics based controller,” in *Advanced Intelligent Mechatronics, 2009. AIM 2009. IEEE/ASME International Conference on*. IEEE, 2009, pp. 1891–1896.
- [33] J. Weingarten, M. Buehler, R. Groff, and D. Koditschek, “Gait generation and optimization for legged robots,” in *The IEEE International Conference on Robotics and Automation*, 2002.
- [34] M. T. Hale, A. M. Johnson, and D. E. Koditschek, “Novel Proprioceptive Sensor for a Legged Robot,” 2011, in preparation.